IAP20 Rec'd PCT/PTO 20 JAN 2006

EFFICIENT MAPPING OF RECONSTRUCTION ALGORITHMS FOR MAGNETIC RESONANCE IMAGING ONTO A RECONFIGURABLE RECONSTRUCTION SYSTEM

The present invention relates to diagnostic medical imaging. It finds particular application in conjunction with the reconstruction of magnetic resonance images and will be described with particular reference thereto.

Heretofore, magnetic resonance imaging scanners have included a main magnet, typically superconducting, which generates a temporally constant magnetic field B_0 through an examination region. A radio frequency coil, such as a whole-body coil, and a transmitter tuned to the resonance frequency of the dipoles to be imaged in the B_0 field have often been used to excite and manipulate these dipoles. Spatial information has been encoded by driving the gradient coils with currents to create magnetic field gradients in addition to the B_0 field across the examination region in various directions. Magnetic resonance signals have been acquired by the same coil, demodulated, filtered and sampled by an RF receiver and finally reconstructed into an image on some dedicated or general-purpose hardware.

Rather than using the same coil to transmit and receive RF pulses, the use of surface or local receive coils has become more and more common recently. These receive coils are often arranged in arrays, in which each coil element produces its own output. Instead of combining the outputs of the coil elements in the analog domain, it has proven advantageous to reconstruct the output from individual coil elements separately. Therefore, each coil element is typically connected with its own RF receiver.

While current scanners claim to have a few receive channels with independent RF receivers, they still have only a single reconstruction unit. The processing of the data from each of the RF receivers is interleaved in time in the reconstruction unit, although it may be performed in parallel to reduce reconstruction times.

Simply multiplying the reconstruction units gives rise to the problem of how to map the processing efficiently onto the individual units. A fixed allocation of reconstruction units to receive channels, for example, makes only poor use of available hardware since varying numbers of coil elements might be employed in practice.

Moreover, the complexity of the reconstruction software generally increases considerably to divide the processing suitably among the reconstruction units.

The present invention provides an improved imaging apparatus and an improved method, which overcome the above-referenced problems and others.

In accordance with one aspect of the present invention, an MRI system is disclosed. A means creates and transmits RF pulses into an examination region to excite and manipulate a spin system to be imaged. A means picks up an MR signal emitted from the examination region. A means demodulates the MR signal and converts the demodulated MR signal into digital data. A means, including a plurality of reconfigurable processing units with dynamically reconfigurable connections, reconstructs the digital data into images.

In accordance with another aspect of the present invention, a method for processing an MR signal is disclosed. RF pulses are created and transmitted into an examination region to excite and manipulate a spin system to be imaged. The MR signal, emitted from the examination region, is picked up. The picked up MR signal is demodulated and converted into digital data. The digital data is reconstructed into images via a plurality of processing units with dynamically reconfigurable connections.

Advantages of the present invention reside, inter alia, in an increased reconstruction speed due to a more efficient utilization of hardware resources, and simpler reconstruction software architecture due to a single general strategy for mapping processing tasks to hardware resources.

The invention may take form in various components and arrangements of components, and in various steps and arrangements of steps. The drawings are only for purposes of illustrating the preferred embodiments and are not be construed as limiting the invention.

FIGURE 1 is a diagrammatic illustration of a magnetic resonance imaging system in accordance with the present invention;

FIGURE 2 is a diagrammatic illustration of a reconfigurable reconstruction system in accordance with the present invention;

FIGURE 3 is a diagrammatic illustration of a possible distribution of processing tasks over four pipeline stages in accordance with the present invention;

FIGURE 4 is a diagrammatic illustration of a possible timing for executing an iterative reconstruction on four processing units per channel in accordance with the present invention;

FIGURES 5A-B depict two alternative techniques for combining images from individual processing channels to create a final combined image in accordance with the present invention;

FIGURE 6A is a diagrammatic illustration of a reconfigurable reconstruction system utilizing six processing channels with one pipeline stage each in accordance with the present invention;

FIGURE 6B is a diagrammatic illustration of a reconfigurable reconstruction system utilizing three processing channels with two pipeline stages in accordance with the present invention;

FIGURE 6C is a diagrammatic illustration of a reconfigurable reconstruction system utilizing two processing channels with three pipeline stages each in accordance with the present invention;

FIGURES 7A-C are diagrammatic illustrations of a reconfigurable reconstruction system built up of boards comprising six embedded processing units each that supports different numbers of processing channels and pipeline stages while utilizing the same total number of processing units, in accordance with the present invention;

FIGURE 8 is a diagrammatic illustration of a reconfigurable reconstruction system built up of a general-purpose hardware, including personal computers or workstations as processing units and a switch as an interconnection.

With reference to FIGURE 1, a magnetic resonance (MR) imaging scanner **10** includes a preferably superconducting main magnet **12**, which includes a solenoid coil in the illustrated embodiment. The main magnet **12** generates a spatially and temporally

constant magnetic field B_0 through an examination region **14** in a bore **16** of the magnet **12**.

Magnetic field gradients across the examination region **14** are generated by gradient coils **18** to spatially encode an MR signal, to spoil the magnetization, and the like. In the preferred embodiment, the gradient coils **18** produce gradients in three orthogonal directions, including a longitudinal or z-direction and transverse or x- and y-directions.

A whole-body coil **20**, preferably a birdcage coil, transmits radiofrequency (RF) signals for exciting and manipulating a spin system to be imaged and may also receive the MR signal.

A plurality of local RF coils **22** is disposed in the bore **16**. The local coils **22** include in the illustrated embodiment a phased-array coil **24**, which includes seven coil elements. Optionally, the phased-array coil may be built into a patient support **26**. In addition, a surface coil array **28** is disposed in the bore **16**. It may include a plurality of surface coils, coils which view different regions of the subject, coils which view a common region of the subject, but have different reception properties, and the like.

To perform measurements, a subject is placed in the magnet's bore **16** with the region of interest in the examination region at or near the magnet's isocenter. A sequence controller **30** controls the gradient amplifiers **32**, which drive the gradient coils to create gradient magnetic fields with appropriate strength, orientation and timing. The sequence controller **30** also controls the radiofrequency transmitter **34** which, with the help of the whole-body coil **20**, sends radiofrequency pulses into the examination region **14** to excite and manipulate the spin system to be imaged.

Magnetic resonance signals are induced in selected receive coils in the examination region **14**. Each of n elements of the local coil arrays **22** is connected with one of n RF receivers $\mathbf{36}_1, \ldots, \mathbf{36}_n$. The whole-body coil **20** is also preferably connected to one additional RF receiver.

The reconfigurable reconstruction system **40** supports up to n independent processing channels $\mathbf{42}_1, \ldots, \mathbf{42}_n$, with each of these channels connected to one of the RF receivers $\mathbf{36}_1, \ldots, \mathbf{36}_n$. The images reconstructed separately by the processing channels are finally combined by the combining unit **44**. The combined images (and optionally the uncombined images) are sent to the host computer **50** for storage and

viewing. The host computer **50**, preferably a personal computer or workstation, includes a display and a user interface connected with the sequence controller **30**, which allows the operator to select among a variety of sequences and imaging parameters.

With continuing reference to FIGURE 1 and further reference to FIGURE 2, the data provided by coils **20, 22, 28** are sent via the RF receivers or receive channels **36_1, ..., 36_n** to corresponding individual channels of a plurality of processing channels **42_1, 42_2, ..., 42_n**. The data are processed by a plurality of processing or reconstruction units **52**, arranged in the pipeline stages **54_1, 54_2, ... 54_m**. The allocation of processing or reconstruction units **52** to processing channels and pipeline stages is performed dynamically on a per scan basis. Moreover, the number of processing channels is adapted to the number of receive channels actually in use, i.e. it is chosen to be a multiple or a factor of the number of active receive channels, or to be the same. The images reconstructed separately by the processing channels **42_1, 42_2, ..., 42_n** are sent to the combining unit **44**, where the images are combined.

With reference to FIGURE 3, one of the processing channels of the reconstruction system is shown in more detail. The reconstruction is performed using four pipeline stages **54_1, 54_2, 54_3**, and **54_4**. The first pipeline stage **54_1** operates on the data in k-space. It performs, for instance, a sampling density compensation or a regridding. The intermediate pipeline stages **54_2** and **54_3** transform the data from k-space to spatial (or image) domain. The use of two pipeline stages permits, in this case, to separate the two-dimensional Fourier transform required in two-dimensional imaging into two subsequent one-dimensional Fourier transforms, allocating one of them to each pipeline stage. The final pipeline stage **54_4** operates on the data in the image domain. It performs, for instance, a roll-off correction or weighting. Alternatively, the images from the individual processing channels are also partly or completely combined in the final pipeline stage to drastically reduce the required bandwidth to the combining unit. In case of an iterative reconstruction, for which a variety of algorithms are known, these processing steps make up the forward processing. Keeping the same mapping of processing tasks to the pipeline stages, the backward processing can be implemented similarly by sending the data in reverse direction from the last to the first pipeline stage. In addition, some further processing in the spatial domain has to be implemented in the last pipeline stage. It includes the core of the iterative reconstruction, such as the conjugate gradient or the

generalized minimum residual method, but without the matrix-vector multiplication, and a redistribution of the final combined image to all processing or reconstruction units allocated to the last pipeline stage before the beginning of a new iteration.

FIGURE 4 shows a possible timing for an iterative reconstruction executed on the four pipeline stages **54_1**, **54_2**, **54_3**, and **54_4** of FIGURE 3. P_xy denotes the processing of image x in iteration y. In the initial iteration, an image A is manipulated in pipeline stages **54_1**, **54_2**, **54_3**, and **54_4** using the forward processing. The images B, C, and D enter pipeline stage **54_1** at suitable later times. When the first image A reaches pipeline stage **54_4**, pipeline stage **54_1** has processed images B, C, and D in the initial iteration. Then, the backward processing starts with the image A in the first iteration on pipeline stage **54_4**. Preferably, a first chain of processors is dedicated to the forward processing and a second chain of processors is dedicated to the backward processing, although the forward and backward processing can also be executed, even simultaneously, on the same processors.

In FIGURE 5A and 5B, exemplary techniques for combining images reconstructed separately by the processing channels are shown. The combination is performed by the processing or reconstruction units allocated to the last pipeline stage **54_m**, which have the capability of exchanging data with each other.

In FIGURE 5A, the image from channel **42_1** is combined with the image from channel **42_2**, producing an intermediate combined image, which is sent to the adjacent channel **42_3** to be further combined with the image from this channel. At the same time, the image from channel **42_n** is combined with the image from channel **42_{n-1}**, producing an intermediate combined image, which is sent to the adjacent channel **42_{n-2}** to be further combined with the image from this channel. After the final combined image from all channels **42_1**, **42_2**, ..., **42_n** has been obtained after n/2 steps, it is sent to the combining unit **44** for further processing.

In FIGURE 5B, the images from channels **42_1** and **42_2**, **42_3** and **42_4**, ..., **42_{n-1}** and **42_n** are combined in parallel. After the final combined image from all channels **42_1**, **42_2**, ..., **42_n** has been obtained, it is sent to the combining unit **44** for further processing. Alternatively, the combination process may be stopped earlier and all remaining intermediate combined images may be sent to the combining unit **44** for further processing.

FIGURES 6A-C illustrate exemplary implementations of the present invention utilizing six processing or reconstruction units **52_1, 52_2, ..., 52_6**. In FIGURE 7, six processing or reconstruction units **52_1, 52_2, ..., 52_6** are configured to process six channels **42_1, 42_2, ..., 42_6**, with a single pipeline stage **54_1** each. The data from six coil elements are sent to six corresponding processing channels. The six images from each of the processing channels are summed up in the combining unit **44**.

In FIGURE 6B, six processing or reconstruction units **52_1, 52_2, ..., 52_6** are configured to process three channels **42_1, 42_2**, and **42_3** with two pipeline stages **54_1, 54_2** each. The data from three coil elements are sent to three corresponding processing channels. The three images from each of the processing channels are summed up in the combining unit **44**.

In FIGURE 6C, six processing or reconstruction units **52_1, 52_2, ..., 52_6** are configured to process two channels **42_1** and **42_2** with three pipeline stages **54_1, 54_2** and **54_3** each. The data from two coil elements are sent to two corresponding processing channels. The two images from each of the processing channels are summed up in the combining unit **44**.

FIGURES 7A-C and 8 show two alternative implementations of the interconnections between the six processing or reconstruction units **52_1, 52_2, ..., 52_6** of FIGURES 6A-C using a switch **60** or other hardware with similar functionality. The interconnections can be configured to realize the network topologies of FIGURES 6A-C. Although six processing units are shown by way of example, any number of processors could be used.

In FIGURE 7A, a crossbar switch **60** is used to connect the six embedded processors **52_1, 52_2, ..., 52_6** of FIGURE 6A, which allows a static configuration of the connections **56** in hardware on a per scan basis. Each processor receives input data separately via the inputs I_1 through I_6. The processors **52_1, 52_2, ..., 52_6** exchange images with each other via the crossbar **60**. After completion of reconstruction, each processor sends an image via the outputs O_1 through O_6 to the combining unit **44**. Alternatively, the image combination is performed partly or entirely on the processors themselves, as discussed above.

In FIGURE 7B, a crossbar switch **60** is used to connect the six embedded processors **52_1, 52_2, ..., 52_6** as shown in FIGURE 6B. The processors **52_1, 52_3**, and **52_5**

are allocated to the pipeline stage 54_1 of channels 42_1, 42_2, and 42_3. The processors 52_1, 52_3, and 52_5 receive input data via the inputs I_1 through I_3. The processors 52_2, 52_4, and 52_6 are allocated to the pipeline stage 54_2 of channels 42_1, 42_2, and 42_3. The processors 52_2, 52_4, and 52_6 exchange images with each other via the crossbar **60**. After completion of reconstruction, the processors 52_2, 52_4, and 52_6 send images via the outputs O_1 through O_3 to the combining unit **44**.

In FIGURE 7C, a crossbar switch **60** is used to connect the six embedded processors 52_1, 52_2, ..., 52_6 as shown in FIGURE 6C. The processors 52_1 and 52_4 are allocated to the pipeline stage 54_1 of channels 42_1 and 42_2. The processors 52_1 and 52_4 receive input data via the inputs I_1 and I_2. The processors 52_3 and 52_6 are allocated to the pipeline stage 54_3 of channels 42_1 and 42_2. The processors 52_3 and 52_6 exchange images with each other via the crossbar **60**. After completion of the reconstruction, the processors 52_3 and 52_6 send images via the outputs O_1 and O_2 to the combining unit **44**.

In FIGURE 8, a switched fabric switch **60** is used to connect the six personal computers or workstations 52_1, 52_2, ..., 52_6, each serving as one processing or reconstruction unit. The switch **60** permits a dynamic configuration of the connections **56** in software for each packet of data.

Thus, the systems shown in FIG. 7A-C and 8 can be configured for a first scan to have six processing channels with one pipeline stage each as per FIG. 7A; for a second scan to have three processing channels with two pipeline stages each as per FIG. 7B; and for a third scan to have two processing channels with three pipeline stages each as per FIG. 7C. Further, each processing or reconstruction unit need not be dedicated to a specific channel. Rather, one or more of the processing or reconstruction units can be shared between two or more channels.

The invention has been described with reference to the preferred embodiments. Modifications and alterations may occur to others upon reading and understanding the preceding detailed description. It is intended that the invention be constructed as including all such modifications and alterations insofar as they come within the scope of the appended claims or the equivalents thereof.